

# Split Bregman Method for Sparse Inverse Covariance Estimation with Matrix Iteration Acceleration

Gui-Bo Ye and Jian-Feng Cai and Xiaohui Xie

*School of Information and Computer Science, University of California, Irvine, CA 92697, USA*  
e-mail: [yeg@uci.edu](mailto:yeg@uci.edu)

*Department of Mathematics, University of California, Los Angeles, CA 90095, USA*  
e-mail: [cai@math.ucla.edu](mailto:cai@math.ucla.edu)

*School of Information and Computer Science, University of California, Irvine, CA 92697, USA*  
e-mail: [xhx@ics.uci.edu](mailto:xhx@ics.uci.edu)

**Abstract:** We consider the problem of estimating the inverse covariance matrix by maximizing the likelihood function with a penalty added to encourage the sparsity of the resulting matrix. We propose a new approach based on the split Bregman method to solve the regularized maximum likelihood estimation problem. We show that our method is significantly faster than the widely used graphical lasso method, which is based on blockwise coordinate descent, on both artificial and real-world data. More importantly, different from the graphical lasso, the split Bregman based method is much more general, and can be applied to a class of regularization terms other than the  $\ell_1$  norm.

**Keywords and phrases:** Inverse covariance matrix, Split Bregman,  $\ell_1$  norm, Graphical model.

## 1. Introduction

Undirected graphical models provide an efficient way to describe and explain the relationships among a set of objects (or variables), and have become a popular tool to model networks of components in a variety of applications, including internet, social networks, and gene networks [1, 2, 3]. A key step involved in constructing a graphical model for a particular application is to learn the structure of the graphical model from a set of observations, which is often called model selection in statistics. If we assume the observations have a multivariate Gaussian distribution with mean  $\mu$  and covariance matrix  $\Sigma$ , learning the structure of a graphical model can be reformulated as a covariance selection problem [4], which aims to identify nonzero entries of the inverse covariance matrix (also known as *concentration matrix* or *precision matrix*). The idea behind this formulation is that if the  $(i, j)$ -entry of  $\Sigma^{-1}$  is zero, then the variable  $i$  and  $j$  are conditionally independent, given the other variables [5, 6].

The principle of parsimony suggests that among the graphical models that adequately explains the data we should select the simplest. Thus, it is natural to impose an  $\ell_1$  penalty for the estimation of  $\Sigma^{-1}$  to promote the sparsity of the resulting graph, as has been proposed by a number of authors [7, 2, 8, 9].

It has been noted that if we fit a linear regression model to each variable using the others as predictors, the regression coefficient of variable  $j$  on  $i$  will be, up to a positive scalar, equal to the  $(i, j)$ -entry of  $\Sigma^{-1}$  [10, 6]. Based on this observation, Meinshausen and Bühlmann [7] proposed a simple approach to the structural learning problem; they estimate a sparse graphical model by fitting a lasso regression model for each variable, treating the variable as the response and the others as predictors. The  $(i, j)$ -entry of  $\Sigma^{-1}$  is estimated to be non-zero if either the estimated coefficient of variable  $i$  on  $j$  or the estimated coefficient of variable  $j$  on  $i$  is nonzero (alternatively, one can use an AND rule, requiring both coefficients to be nonzero). Although this approach is computationally attractive and has been shown to be able to consistently estimate  $\Sigma^{-1}$  asymptotically, it does not take the intrinsic symmetry of  $\Sigma^{-1}$  into account, and could result in contradictory neighborhoods and a non-positive definite concentration matrix. Furthermore, if the same penalty parameter is used for all lasso regressions, as is commonly done to reduce the number of parameters, the approach will not be able to correctly infer graphs with skewed degree distributions. To overcome these limitations, Peng et al. [6] proposed a symmetric regression approach; however, the derived concentration matrix can still be non-positive-definite.

A more principled approach to covariance selection is to find the  $\Sigma^{-1}$  that maximizes the log-likelihood of the data with an added  $\ell_1$  penalty [11] to encourage the sparsity of the resulting graph [2, 8, 9]. More

specifically, suppose we are given  $n$  samples independently drawn from a  $p$ -variate Gaussian distribution:  $\mathbf{x}_1, \dots, \mathbf{x}_n \sim \mathcal{N}(\mu, \Sigma)$ . Let  $S$  be the empirical covariance matrix:

$$S := \frac{1}{n} \sum_{k=1}^n (\mathbf{x}_i - \mu)(\mathbf{x}_i - \mu)^T.$$

Denote  $\Theta = \Sigma^{-1}$ . The goal is to find the  $\Theta^*$  that maximizes the penalized log-likelihood

$$\log \det \Theta - \text{tr}(S\Theta) - \lambda \sum_{i \neq j} |\Theta_{ij}|, \quad (1)$$

subject to the constraint that  $\Theta$  is positive definite. Here, we use  $\text{tr}$  to denote the trace and  $\Theta_{ij}$  to denote the  $(i, j)$ -entry of  $\Theta$ . Due to the  $\ell_1$  penalty term and the explicit positive definite constraint on  $\Theta$ , the method leads to a sparse estimation of the concentration matrix that is guaranteed to be positive definite. The simpler approach of Meinshausen and Bühlmann [7] can be viewed as an approximation to the penalized maximum likelihood formulation [8, 9].

The objective function in (1) is strictly convex, so a global optimal solution is guaranteed to exist and be unique. Finding the optimal solution can, however, be computationally challenging due to the log det term appeared in the likelihood function and the nondifferentiability of the  $\ell_1$  penalty. Yuan and Lin [2] solved (1) using the interior-point method for the “maxdet” problem, which is prohibitive for problems with more than tens of variables due to its memory requirements and computational complexity. Banerjee et al. [8] developed a blockwise coordinate descent method to solve (1), after noting that the dual of (1) can be reduced to a lasso regression problem if one focuses on optimizing only one row or column of  $\Omega$ . Exploiting this observation further, Friedman et al. [9] used the coordinate descent procedure to solve the lasso regression arising in the dual of the blockwise coordinate descent, and implemented an efficient method, called “graphical lasso”, to solve (1). The graphical lasso is remarkably fast; it can solve a  $p = 1000$  dimension ( $\sim 500,000$  parameters) problem within a minute. However, a major limitation of the blockwise coordinate descent methods, including the graphical lasso, is that their derivation is specific to the  $\ell_1$  penalty. This limitation prevents them from being applied to other types of penalties, such as a combination of  $\ell_1$  and  $\ell_2$  penalty used in the elastic net [12], the fused lasso penalty [13], and the group lasso penalty [14], which have been found to be more useful than the simple  $\ell_1$  penalty in a number of applications.

In this paper, we propose a new algorithm based on the split Bregman method [15, 16] to solve the penalized maximum likelihood estimation problem (1). We show that our method is not only substantially faster than the graphical lasso method, but can also be easily extended to deal with a broad range of penalty terms.

Although the Bregman iteration was an old technique proposed in the sixties [17, 18], it gained significant interest only recently after Osher and others demonstrated its high efficiency for sparsity recovery in a wide range of problems, including image restoration [19, 15, 16], compressed sensing [20, 21, 22], matrix completion [23], low rank matrix recovery [24], and general fused lasso problems [25]. By introducing an auxiliary variable, we show that the optimization problem in (1) can be reformulated such that the log-likelihood term and the penalty term interact only through an equality constraint, thereby enabling an efficient application of the split Bregman method to solve the optimization problem.

The rest of the paper is organized as follows. In Section 2, we first present a generalized sparse inverse covariance matrix estimation problem, which includes (1) as a special case. We then derive a split Bregman algorithm, called SBGM, to solve the generalized problem (Subsection 2.2). The convergence property of the algorithm is also given. SBGM consists of three update steps with the first update being the major time-consuming one. In Subsection 2.3, we provide an explicit formula for the first update step and propose an efficient Newton method to solve the resulting matrix quadratic equation. In Section 3, we implement SBGM to solve the special case (1). In Section 4, we illustrate the utility of our algorithm and compare its performance to the graphical lasso method using both simulated data and gene expression data.

## 2. The split Bregman method for generalized sparse graphical models (SBGM)

The split Bregman method was originally proposed by Osher and coauthors to solve total variation based image restoration problems [15]. It was later found to be either equivalent or closely related to a number of

other existing optimization algorithms, including Douglas-Rachford splitting [26], the alternating direction method of multipliers (ADMM) [27, 28, 15, 29] and the method of multipliers [30]. Because of its fast convergence and the easiness of implementation, it is increasingly becoming a method of choice for solving large-scale sparsity recovery problems [24, 16, 25].

In this section, we first extend the  $\ell_1$  regularized maximum likelihood inverse covariance matrix estimation problem (1) to allow for a more general class of regularization terms. We reformulate the problem by introducing an auxiliary variable. We then proceed to derive a split Bregman method to solve the reformulated problem.

### 2.1. A general problem and its reformulation

We derive our algorithm in a more general setting than the one described in (1). Instead of using the  $\ell_1$ -norm penalty, we consider a more general penalty  $\phi(\Theta)$ , which is only required to be convex and satisfy  $\phi(\Theta) = \phi(\Theta^T)$ . The  $\ell_1$  norm and a number of other types, including those used in elastic net [12], fused lasso [13], and group lasso [14], can be viewed as a special case of the general penalty term. We find  $\Theta^*$  by solving the following constrained optimization problem

$$\min_{\Theta \succ 0} -\log \det \Theta + \text{tr}(S\Theta) + \lambda\phi(\Theta), \quad (2)$$

where  $\Theta \succ 0$  means that  $\Theta$  is positive definite, and  $\lambda \in \mathbb{R}_+$  is a regularization parameter. If we choose  $\phi(\Theta) = \sum_{i \neq j} |\Theta_{ij}|$ , the general problem reduces to the problem (1).

The log-likelihood term and the regularization term in (2) are coupled, which makes the optimization problem difficult to solve. However, the two terms can be decoupled if we introduce an auxiliary variable to transfer the coupling from the objective function to the constraints. More specially, the problem (2) is equivalent to the following problem

$$\begin{aligned} \min \quad & -\log \det \Theta + \text{tr}(S\Theta) + \lambda\phi(A) \\ \text{s.t.} \quad & A = \Theta \\ & \Theta \succ 0. \end{aligned} \quad (3)$$

The introduction of the new variable of  $A$  is a key step of our algorithm, which makes the problem amenable to a split Bregman procedure to be detailed below.

### 2.2. Derivation of the split Bregman method

Although the split Bregman method originated from Bregman iterations [31, 16, 32], it has been demonstrated to be equivalent to the alternating direction method of multipliers (ADMM) [27, 28, 33, 34]. For simplicity of presentation, next we derive the split Bregman method using the augmented Lagrangian method [35, 30].

We first define an augmented Lagrangian function of (3)

$$\mathcal{L}(\Theta, A, M) = -\log \det \Theta + \text{tr}(S\Theta) + \lambda\phi(A) + \text{tr}(M^T(\Theta - A)) + \frac{\mu}{2}\|\Theta - A\|_F^2, \quad (4)$$

where matrix  $M$  is a dual variable corresponding to the linear constraint  $\Theta = A$ , and  $\mu > 0$  is a parameter. Compared with the standard Lagrangian function, the augmented Lagrangian function has an extra term  $\frac{\mu}{2}\|\Theta - A\|_F^2$ , which penalizes the violation of the linear constraint  $\Theta = A$ .

With the definition of the augmented Lagrangian function (4), the primal problem (3) is equivalent to

$$\min_{\Theta \succ 0, A} \max_M \mathcal{L}(\Theta, A, M). \quad (5)$$

Exchanging the order of min and max in (5) leads to the formulation of the dual problem

$$\max_M E(M), \quad \text{with} \quad E(M) = \min_{\Theta \succ 0, A} \mathcal{L}(\Theta, A, M). \quad (6)$$

Note that the gradient  $\nabla E(M)$  can be calculated by the following [36]

$$\nabla E(M) = \Theta(M) - A(M), \quad \text{with} \quad (\Theta(M), A(M)) = \arg \min_{\Theta \succ 0, A} \mathcal{L}(\Theta, A, M). \quad (7)$$

Applying gradient ascent on the dual problem (6) and using equation (7), we obtain the method of multipliers [30] to solve (3)

$$\begin{cases} (\Theta^{k+1}, A^{k+1}) = \arg \min_{\Theta \succ 0, A} \mathcal{L}(\Theta, A, M^k), \\ M^{k+1} = M^k + \mu(\Theta^{k+1} - A^{k+1}). \end{cases} \quad (8)$$

Here we have used  $\mu$  as the step size of the gradient ascent. It is easy to see that the efficiency of the iterative algorithm (8) largely hinges on whether the first equation of (8) can be solved efficiently. Note that the augmented Lagrangian function  $\mathcal{L}(\Theta, A, M^k)$  still contains a nondifferentiable term  $\phi(A)$ . But different from the original objective function (2), the function  $\phi$  induced nondifferentiable term has now been transferred from terms involving  $\Theta$  to terms involving  $A$  only. Thus we can solve the first equation of (8) through an iterative algorithm that alternates between the minimization of  $\Theta$  and  $A$ ,

$$\begin{cases} \Theta^{k+1} = \arg \min_{\Theta \succ 0} -\log \det \Theta + \text{tr}(S\Theta) + \text{tr}((M^k)^T(\Theta - A^k)) + \frac{\mu}{2} \|\Theta - A^k\|_F^2, \\ A^{k+1} = \arg \min_A \lambda \phi(A) + \text{tr}((M^k)^T(\Theta^{k+1} - A)) + \frac{\mu}{2} \|\Theta^{k+1} - A\|_F^2. \end{cases} \quad (9)$$

The method of multipliers requires that the alternative minimization of  $\Theta$  and  $A$  in (9) be run multiple times until convergence. However, because the first equation of (8) represents only one step of the overall iteration, it is actually not necessary to be solved completely. In fact, the split Bregman method (or the alternating direction method of multipliers [27]) uses only one iteration of (8), which leads to the following iterative algorithm for solving (3),

$$\begin{cases} \Theta^{k+1} = \arg \min_{\Theta \succ 0} -\log \det \Theta + \text{tr}(S\Theta) + \text{tr}((M^k)^T(\Theta - A^k)) + \frac{\mu}{2} \|\Theta - A^k\|_F^2, \\ A^{k+1} = \arg \min_A \lambda \phi(A) + \frac{\mu}{2} \|\Theta^{k+1} - A + \mu^{-1} M^k\|_F^2, \\ M^{k+1} = M^k + \mu(\Theta^{k+1} - A^{k+1}). \end{cases} \quad (10)$$

### 2.2.1. Convergence Property

The convergence of the iteration (10) can be derived from the convergence theory of the alternating direction method of multipliers or the convergence theory of the split Bregman method [27, 37, 16].

**Theorem 1.** *Let  $\Theta^k$  be generated by (10), and  $\Theta^*$  be the unique minimizer of (3). Then,*

$$\lim_{k \rightarrow \infty} \|\Theta^k - \Theta^*\| = 0.$$

From Theorem 1, the condition for the convergence of the iteration (10) is quite mild and even irrelevant to the choice of the parameter  $\mu$  in the iteration (10). This property makes the split Bregman method quite general and easy to implement, which partly explains why the method is gaining popularity recently.

### 2.2.2. Updating $\Theta$ and $A$

We first focus on the computation of the first equation of (10). Taking the derivative of the objective function and setting it to be zero, we get

$$-\Theta^{-1} + \mu\Theta = \mu A^k - S - M^k. \quad (11)$$

It is a quadratic equation where the unknown is a matrix. The complexity for solving this equation is at least  $O(p^3)$  because of the inversion involved in (11). Note that because  $\phi(\Theta) = \phi(\Theta^T)$ , if  $\Theta^k$  is symmetric,

so is  $\mu A^k - S - M^k$ . It is easy to check that the explicit form for the solution of (11) under constraint  $\Theta \succ 0$ , i.e.,  $\Theta^{k+1}$ , is

$$\Theta^{k+1} = \frac{K^k + \sqrt{(K^k)^2 + 4\mu I}}{2\mu}, \quad \text{where } K^k = \mu A^k - S - M^k. \quad (12)$$

Here  $\sqrt{C}$  is the square root of a symmetric positive definite matrix  $C$ . Recall that the square root of a symmetric positive definite matrix  $C$  is defined to be the matrix whose eigenvectors are the same as those of  $C$  and eigenvalues are the square root of those of  $C$ . Therefore, to get the update of  $\Theta^{k+1}$ , one may first compute the eigenvalues and eigenvectors of  $K^k$ , and then get the eigenvalues of  $\Theta^{k+1}$  according to (12) by replacing the matrices by the corresponding eigenvalues. This approach is rather slow due to the eigen decomposition step. In the next subsection, we will propose a faster algorithm for solving (11) based on Newton's iteration [38, 39].

For the second equation of (10), we have made the data fitting term  $\frac{\mu}{2} \|\Theta^{k+1} - A + \mu^{-1} M^k\|_F^2$  separable with respect to the entries of  $A$ . Thus, if the  $\phi(A)$  is separable, then it is very easy to get the solution and the computational complexity would be  $O(p^2)$  for most cases. See Section 3 for the special case of  $\phi(A) = \sum_{i \neq j} |A_{ij}|$ . Thus, compared with the update of  $\Theta$ , the computational time for updating  $A$  is minor and mostly negligible. The computation of the third equation of (10) is straightforward and the computational cost is also  $O(p^2)$ , which can be neglected as well. So the overall computational time for running one iteration of (10) mostly depends on how fast  $\Theta$  can be updated.

### 2.3. Newton's method to calculate the square root of a positive definite matrix

As described above, the main step to obtain an update of  $\Theta^{k+1}$  is calculating the square root of the positive definite matrix  $(K^k)^2 + 4\mu I$ , where  $K^k$  is a symmetric matrix. One possible way is to calculate all the eigenvalues and eigenvectors of matrix  $K^k$ , which is computational demanding when the size of the matrix is large. In this subsection, we will use Newton's method [39, 38] to calculate the square root of a positive definite matrix directly instead of using the eigenvalue decomposition of  $K^k$ . Our numerical experiments show that Newton's method for calculating the square root of a positive definite matrix of the form  $K^2 + 4\mu I$ , where  $K$  is symmetric, is about four times faster than the one using eigen decomposition.

We begin with finding the positive root of the equation  $x^2 - a = 0 (a > 0)$  using Newton's method [40, 41]. That is,

$$x^{k+1} = \frac{1}{2} \left( x^k + \frac{a}{x^k} \right). \quad (13)$$

The following lemma ensures that (13) with  $x^0 > 0$  always converges to  $\sqrt{a}$  and the convergence rate is quadratic.

**Lemma 1.** *If we choose  $x^0 > 0$ , then  $x^k$  generated by (13) is well-defined and satisfies*

$$|x^{k+1} - \sqrt{a}| \leq \min \left\{ \frac{1}{2\sqrt{a}} |x^k - \sqrt{a}|^2, \frac{1}{2} |x^k - \sqrt{a}| \right\}.$$

*Proof.* Since  $x^0 > 0$ ,  $x^1$  is well-defined and  $x^1 = \frac{1}{2} \left( x^0 + \frac{a}{x^0} \right) \geq \sqrt{a}$ . By induction,  $x^k$  is well-defined and  $x^k \geq \sqrt{a}$ . Moreover, using the iteration (13),

$$|x^{k+1} - \sqrt{a}| = \left| \frac{1}{2} \left( x^k + \frac{a}{x^k} \right) - \sqrt{a} \right| = \frac{1}{2|x^k|} |x^k - \sqrt{a}|^2 \leq \frac{1}{2\sqrt{a}} |x^k - \sqrt{a}|^2,$$

and

$$|x^{k+1} - \sqrt{a}| = \frac{1}{2|x^k|} |x^k - \sqrt{a}|^2 = \frac{1}{2} \left( 1 - \frac{\sqrt{a}}{x^k} \right) |x^k - \sqrt{a}| \leq \frac{1}{2} |x^k - \sqrt{a}|.$$

□

Now we apply Newton's method to calculate the square root of a positive definite matrix of the form  $K^2 + \alpha I$ , where  $K$  is symmetric and  $\alpha$  is a positive constant. Let  $X^0 = \sqrt{\alpha}I$  and

$$X^{k+1} = \frac{1}{2} (X^k + (X^k)^{-1}(K^2 + \alpha I)). \quad (14)$$

Then the iteration (14) always converges quadratically to  $\sqrt{K^2 + \alpha I}$ . Similar algorithms have been proposed in [42, 43, 41] to compute polar decompositions.

**Theorem 2.** *If we choose  $X^0 = \sqrt{\alpha}I$ , then  $X^k$  generated by (14) is well-defined and quadratically converges to  $\sqrt{K^2 + \alpha I}$ . More specifically,*

$$\|X^k - \sqrt{K^2 + \alpha I}\|_2 \leq \min \left\{ \frac{1}{2\sqrt{\alpha + \lambda_{\min}(K^2)}} \|X^k - \sqrt{K^2 + \alpha I}\|_2^2, \frac{1}{2} \|X^k - \sqrt{K^2 + \alpha I}\|_2 \right\}, \quad (15)$$

where  $\lambda_{\min}(K^2)$  is the minimum eigenvalue of matrix  $K$ .

*Proof.* Let  $K = U\Omega U^T$  with  $\Omega = \text{diag}(\omega_1, \dots, \omega_p)$  and  $UU^T = U^T U = I$ . Since  $X^0 = \sqrt{\alpha}I$ , by induction,  $X^k$  can be written as  $X^k = U\Omega^k U^T$  with  $\Omega^k = \text{diag}(\omega_1^k, \dots, \omega_p^k)$  and  $\omega_i^k > 0$ . Furthermore,

$$X^{k+1} = U \left( \frac{1}{2} (\Omega^k + (\Omega^k)^{-1}(\Omega^2 + \alpha I)) \right) U^T.$$

Therefore, (14) changes only the singular values which are governed by (13). The Theorem follows immediately from Lemma 1.  $\square$

Note that the choice of the initial guess  $X^0$  can be any matrix of the form  $U\Lambda U^T$ , where  $U$  is the eigenvectors of  $K$  and  $\Lambda$  is a diagonal matrix with positive diagonal. In the above, we have chosen  $X^0 = \sqrt{\alpha}I$  for simplicity.

Combining the iteration (10) and Newton's method to calculate the square root of the positive matrix of the form  $K^2 + 4\mu I$  where  $K$  is symmetric, we get Algorithm 1 to compute the solution of (2).

---

**Algorithm 1** Split Bregman method for generalized graphical models (2) (SBGM)

---

Given  $\mu$ . Initialize  $\Theta^0, A^0$ , and  $M^0$ .

**repeat**

- 1) Compute  $K^k = \mu A^k - S - M^k$
- 2) Use Newton method to compute  $X^k = \sqrt{(K^k)^2 + 4\mu I}$
- 3)  $\Theta^{k+1} = \frac{K^k + X^k}{2\mu}$
- 4)  $A^{k+1} = \arg \min_A \lambda \phi(A) + \frac{\mu}{2} \|\Theta - A^k + \mu^{-1} M^k\|_F^2$
- 5)  $M^{k+1} = M^k + \mu(\Theta^{k+1} - A^{k+1})$

**until**

Convergence

---

### 3. The Split Bregman method for $\ell_1$ penalized graphical models (SBGLasso)

In this section, we describe a detailed implementation of the split Bregman method to solve the  $\ell_1$  penalized inverse covariance matrix estimation problem (1). The  $\ell_1$  penalty corresponds to a special case of the general penalty used in (2) with  $\phi(\Theta) = \sum_{i \neq j} |\Theta_{ij}|$ .

Algorithm 1 can be applied here with only minor changes. The updates of  $\Theta$  and  $M$  are exactly the same as the ones in Algorithm 1. For the update of  $A$ , let  $\mathcal{T}_\lambda$  be a soft thresholding operator defined on matrix space and satisfying

$$\mathcal{T}_\lambda(\Omega) = (t_\lambda(\omega_{ij}))_{\substack{i,j=1 \\ i \neq j}}^p, \quad \text{with} \quad t_\lambda(\omega_{ij}) = \text{sgn}(\omega_{ij}) \max\{0, |\omega_{ij}| - \lambda\}.$$

Then the update of  $A$  is

$$A^{k+1} = \mathcal{T}_\mu(\Theta^{k+1} + \mu^{-1} M^k).$$

So we obtain Algorithm 2 to solve (1).

---

**Algorithm 2** Split Bregman method for  $\ell_1$  penalized graphical models (SBGLasso)

---

Initialize  $\Theta^0, A^0$ , and  $M^0$ .  
**repeat**  
  1) Compute  $K^k = \mu A^k - S - M^k$   
  2) Use Newton method to compute  $X^k = \sqrt{(K^k)^2 + 4\mu I}$   
  3)  $\Theta^{k+1} = \frac{K^k + X^k}{2\mu}$   
  4)  $A^{k+1} = \mathcal{T}_{\Delta}(\Theta^{k+1} + \mu^{-1}M^k)$   
  5)  $M^{k+1} = M^k + \mu(\Theta^{k+1} - A^{k+1})$   
**until**  
  Convergence

---

#### 4. Numerical experiments

In this section, we use time trials to illustrate the utility of our proposed algorithms. We first demonstrate the efficiency of the proposed Newton’s method for calculating the square root of a positive definite matrix by comparing it with the eigenvalue decomposition method. We then benchmark the performance of the split Bregman method (SBGLasso) for solving the  $\ell_1$  penalized inverse covariance matrix estimation problem, and compare it to the graphical lasso method. The time trials were all conducted on an Intel Core 2 Duo desktop PC (E7500, 2.93GHz).

##### 4.1. Newton’s method versus eigenvalue decomposition method for computing the square root of a positive definite matrix

We first illustrate the efficiency of the Newton’s method for computing the square root of a positive definite matrix of form  $M = K^2 + \mu I$ , where  $K = (B + B')/2$  and  $B$  is a  $p \times p$  random Gaussian matrix with its entries randomly drawn from the standard Gaussian distribution. The form of the matrix  $M$  we choose is mostly motivated by the square root we need to solve in the equation (12). Our algorithm is implemented in Matlab using mex programming. We compare it with the method using matlab function “schur” on the matrix  $K$ . More specifically, we use the two-line code “[ $U, D$ ] = schur( $K$ );  $SR = U * \text{diag}(\text{sqrt}(\text{diag}(D).^2 + \mu)) * U'$ ” to compute the square root of  $M$ . Note that the upper triangular matrix produced by schur decomposition is exactly diagonal since  $K$  is symmetric.

In Table 1, we have listed the number of iterations needed for our Newton’s method to solve the square root of  $M$ , where we stop the iterations whenever the relative change of two successive steps is less than  $10^{-6}$ . We also calculated the relative difference between the result derived by our method and the one by schur decomposition. As shown in Table 1, the relative difference is of precision of order  $10^{-9}$ , demonstrating that our algorithm is highly accurate. Empirically, we find that our algorithm only needs 8 steps to converge and the number of steps is mostly constant as the matrix size increases, which illustrates the quadratic convergence rate of our algorithm given in Theorem 2. Overall, our algorithm is substantially faster than the schur decomposition based method, saving more than 75% of the computational times on all examples we tested. For example, when  $n = 2000$ , our algorithm takes 7.45 seconds to find a solution, significantly lower than the 38.08 seconds used by the method based on schur decomposition.

TABLE 1  
*Experimental results for computing the square root of  $M$ . All results are averages of 10 runs.*

p	Newton method		Schur decomposition	relative difference
	iters in Newton method	total time(s)	total time(s)	
1000	8	1.20	4.54	$9.69 \times 10^{-14}$
1500	8	3.35	16.11	$8.54 \times 10^{-13}$
2000	8	7.45	38.08	$3.09 \times 10^{-11}$
2500	8	13.56	76.03	$3.75 \times 10^{-10}$
3000	8	22.43	127.32	$2.35 \times 10^{-9}$



## 4.2. SBGLasso versus the graphical lasso for $\ell_1$ penalized graphical models

Next we illustrate the efficiency of the split Bregman method for  $\ell_1$  penalized maximum likelihood graphical models (SBGLasso) using time trials. The graphical lasso proposed by Friedman et al. [9] is by far the most efficient method for solving the large-scale inverse covariance matrix estimation problem (1), so we focus our comparison with the graphical lasso method. SBGLasso is coded in C, linked to a Matlab function, while the graphical lasso is coded in Fortran, linked to an R language function, so it is reasonable to compare them using time trials despite that they are implemented in two different languages.

The stopping criteria of SBGLasso is specified as follows. Let  $\Phi(\Theta^k) = -\log \det \Theta^k + \text{tr}(S\Theta^k) + \lambda \sum_{i \neq j} |\Theta_{ij}^k|$ . Since the  $\ell_1$  penalized maximum likelihood graphical model (1) is a special case of model (2), we have  $\lim_{k \rightarrow \infty} \Phi(\Theta^k) = \Phi(\Theta^*)$  by Theorem 1. We terminate the algorithm when the relative change of the energy functional  $\Phi(\Theta)$  falls below a certain threshold  $\delta$ . In addition, we require the relative difference  $\|\Theta - A\|_F / \|\Theta\|_F$  to be less than  $\delta$  to make sure that the resulting solution is also primal feasible. We used  $\delta = 10^{-4}$  in our simulation. For the graphical lasso, we also set the termination threshold to be  $10^{-4}$ .

Note that the convergence of Algorithm 2 is guaranteed no matter what values of  $\mu$  are chosen as shown in Theorem 1. However, the choice of  $\mu$  can influence the speed of the algorithm since it could affect the number of iterations involved. In our implementation, we found empirically that choosing  $\mu = 0.5$  works well for the artificial data and  $\mu = 0.012$  for the gene expression data.

### 4.2.1. Artificial data

To generate the artificial data, we first create a set of sparse inverse covariance matrices with different dimension  $p$ , and then generate  $n$  samples from the multivariate Gaussian distribution parametrized by each of the inverse covariance matrices. The sparse inverse covariance matrices are created using the same procedure as described in [8]. More specially, to create an inverse covariance matrix, we first generate a random  $p \times p$  diagonal matrix with positive entries, and then symmetrically insert random numbers to approximately  $p$  randomly chosen locations of the matrix. Positive definiteness is ensured by adding a multiple of the identity to the matrix if needed. The penalty parameter is chosen such that the estimated inverse covariance matrix has roughly the same number of nonzero entries as the actual inverse covariance matrix.


Table 2 shows a comparison of SBGLasso and the graphical lasso on both computational times as well as relative errors after testing on the artificial data. We note that SBGLasso is significantly faster than the graphical lasso for large-scale data ( $p > 500$ ), while being able to achieve the same levels of accuracy. It takes less than half of the computational times used by the graphical lasso for all cases we tested except when  $p = 500$ . For example, when  $n = 2000, p = 3000$ , SBGLasso takes only 269.53 seconds to find the optimal solution, compared to the 729.63 seconds used by the graphical lasso.

To evaluate how the performance of SBGLasso scales with the problem size, we plotted the CPU time that SBGLasso took to solve the problem (1) as a function of  $p$  and  $n$ . The CPU time is roughly quadratic in  $p$  and constant in  $n$ . This phenomena is expected since each of the concentration matrices used in the artificial data has about  $\frac{p(p+1)}{2}$  unknowns, which is quadratic with respect to  $p$ . By contrast, the number of samples only appears in  $S$ , and does not directly affect the speed of SBGLasso.

TABLE 2  
Comparing the performance of SBGLasso and the graphical lasso on the artificial data

$n, p$	SBFGLasso		Graphical Lasso	
	total time (s)	relative error	total time(s)	relative error
$n = 1000, p = 500$	2.61	0.1136	1.94	0.1136
$n = 1000, p = 1000$	14.09	0.1192	27.08	0.1193
$n = 1000, p = 2000$	88.71	0.1170	217.88	0.1170
$n = 2000, p = 3000$	269.53	0.0980	729.63	0.0979





**Fig 1:** CPU times for SBGLasso for the same problem as in Table 2, for different values of  $n$  and  $p$ . (a)  $n$  is fixed and equals to 1000; (b)  $p$  is fixed and equals to 2000.

#### 4.2.2. Gene expression data

We next apply SBGLasso to learn gene regulatory networks from microarray gene expression data. We used the Rosetta Inpharmatics compendium of gene expression profiles described by Hughes et al. [44]. The compendium contains microarray measurements of the expressions of 6316 genes in response to 300 diverse mutations and chemical treatments in *S. cerevisiae*. We preprocessed this data by removing genes that either contain missing values in some experiments, or have low variance across different experimental conditions. This led us to a final dataset of  $p = 1000$  genes with measurements across  $n = 300$  experiments.

We used SBGLasso to solve the sparse graphical model (1) on this data, and compared its performance to the graphical lasso [9]. The results are summarized in Table 3, which shows the computational times spent by different solvers for different choices of the regularization parameter  $\lambda$ . SBGLasso consistently outperforms the graphical lasso, saving more than 40% time in all cases. The concentration matrices derived by SBGLasso and the graphical lasso are very similar. Table 3 also lists the values of the objective function achieved when the algorithms converge, showing that the results from the two solvers are almost identical.

TABLE 3  
Comparing the performance of SBGLasso and the graphical lasso on the gene expression data

$\lambda$	SBFGLasso		Graphical Lasso	
	total time (s)	energy value $\Phi(\Theta)$	total time(s)	energy value $\Phi(\Theta)$
$\lambda = 0.01$	89.10	$-2.96 \times 10^3$	167.23	$-2.96 \times 10^3$
$\lambda = 0.015$	59.95	$-2.78 \times 10^3$	147.92	$-2.79 \times 10^3$
$\lambda = 0.02$	54.41	$-2.72 \times 10^3$	116.18	$-2.72 \times 10^3$
$\lambda = 0.025$	52.94	$-2.69 \times 10^3$	87.74	$-2.69 \times 10^3$

## 5. Discussion

The problem of inverse covariance matrix estimation arises in many applications. Because the number of available samples is usually much smaller than the number of free parameters associated with the inverse covariance matrix, a parsimony approach is to select the simplest matrix that adequately explains the data. This leads to the idea of formulating the problem as a regularized maximum likelihood estimation problem with a penalty added to encourage the sparsity of the resulting matrix. Solving the regularized maximum likelihood problem is however nontrivial for large-scale problems, because of the complexity of the

log-likelihood term and the nondifferentiability of the regularization term. In this work, we propose a new approach based on the split Bregman method to solve the regularized inverse covariance matrix estimation problem. We show that the approach is very fast; it solves a  $3000 \times 3000$  matrix estimation problem in less than 5 minutes.

We compared the split Bregman method to the graphical lasso method, the state-of-the-art for solving  $\ell_1$  penalized inverse covariance matrix estimation problems [9]. We show that our method is about twice faster than the graphical lasso for large-scale problems on both the artificial data and the gene expression data, while being able to achieve the same levels of accuracy. More importantly, our method is much more general and can handle a variety of sparsity regularization terms other than the  $\ell_1$  norm, such as those used in elastic net, fused lasso and grouped lasso. By contrast, the graphical lasso can only be applied to the  $\ell_1$  norm penalty.

The contribution of the paper lies in two aspects. First, we reformulated the regularized maximum likelihood estimation problem by introducing an auxiliary variable, which leads to the decoupling of the log-likelihood term and the regularization term, making the problem amenable to the split Bregman method. Second, we proposed to use Newton’s method to calculate the square root of a positive definite matrix, the major time consuming step of our algorithm. Because the matrix quadratic equation is separable in its eigenvalues, the conventional solver is first to compute the eigen decomposition and then solve univariate quadratic equations about the eigenvalues. This approach is generally slow, as the eigen decomposition is time-consuming. To accelerate it, we propose to use a matrix iterative algorithm [39, 38] to directly solve the matrix quadratic equation based on Newton’s method, without using the eigen decomposition. The contribution is crucial and makes the derived algorithm from the split Bregman method very easy to implement and highly efficient.

## References

- [1] S.L. Lauritzen. *Graphical models*. Oxford University Press, USA, 1996.
- [2] M. Yuan and Y. Lin. Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.
- [3] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.
- [4] A.P. Dempster. Covariance selection. *Biometrics*, 28(1):157–175, 1972. ISSN 0006-341X.
- [5] D. Edwards. *Introduction to graphical modelling*. Springer Texts in Statistics. Springer-Verlag, New York, second edition, 2000.
- [6] J. Peng, P. Wang, N. Zhou, and J. Zhu. Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association*, 104(486):735–746, 2009. ISSN 0162-1459.
- [7] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006. ISSN 0090-5364.
- [8] O. Banerjee, L. El Ghaoui, and A. d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine Learning Research*, 9: 485–516, 2008. ISSN 1532-4435.
- [9] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008. ISSN 1465-4644.
- [10] T. Hastie, R. Tibshirani, and J.H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Verlag, 2009.
- [11] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1): 267–288, 1996.
- [12] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320, 2005.
- [13] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(1):91–108, 2005.
- [14] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67, 2006.
- [15] T. Goldstein and S. Osher. The split Bregman method for  $L_1$ -regularized problems. *SIAM J. Imaging Sci.*, 2(2):323–343, 2009. ISSN 1936-4954.

- [16] J. F. Cai, S. Osher, and Z. Shen. Split bregman methods and frame based image restoration. *Multiscale Model. Simul.*, 8(2):337–369, 2009.
- [17] L. M. Brègman. A relaxation method of finding a common point of convex sets and its application to the solution of problems in convex programming. *Ž. Vyčisl. Mat. i Mat. Fiz.*, 7:620–631, 1967. ISSN 0044-4669.
- [18] A. E. Çetin. Reconstruction of signals from Fourier transform samples. *Signal Process.*, 16(2):129–148, 1989. ISSN 0165-1684.
- [19] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.*, 4(2):460–489 (electronic), 2005. ISSN 1540-3459.
- [20] J. F. Cai, S. Osher, and Z. Shen. Linearized Bregman iterations for compressed sensing. *Math. Comp.*, 78(267):1515–1536, 2009.
- [21] S. Osher, Y. Mao, B. Dong, and W. Yin. Fast linearized Bregman iteration for compressive sensing and sparse denoising. *commu. Math. Sci.*, 8(2):93–111, 2010.
- [22] W. Yin, S. Osher, D. Goldfarb, and J. Darbon. Bregman iterative algorithms for  $l_1$ -minimization with applications to compressed sensing. *SIAM J. Imaging Sci.*, 1(1):143–168, 2008.
- [23] J. F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM J. Optim.*, 20(4):1956–1982, 2010. ISSN 1052-6234.
- [24] E.J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Arxiv preprint arXiv:0912.3599*, 2009.
- [25] G.B. Ye and X. Xie. Split bregman method for large scale fused lasso. *Computational Statistics and Data Analysis*, 2010. .
- [26] C. Wu and X.C. Tai. Augmented Lagrangian method, dual methods, and split Bregman iteration for ROF, vectorial TV, and high order models. *SIAM Journal on Imaging Sciences*, 3:300, 2010.
- [27] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers and Mathematics with Applications*, 2(1):17–40, 1976.
- [28] R. Glowinski and A. Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité, d’une classe de problèmes de dirichlet non linéaires. *Rev. Franc. Automat. Inform. Rech. Operat.*, pages 41–76, 1975.
- [29] E. Esser. Applications of lagrangian-based alternating direction methods and connections to split bregman. *CAM report*, 9:31, 2009. URL <ftp://ftp.math.ucla.edu/pub/camreport/cam09-31.pdf>.
- [30] R. T. Rockafellar. A dual approach to solving nonlinear programming problems by unconstrained optimization. *Math. Programming*, 5:354–373, 1973. ISSN 0025-5610.
- [31] J. F. Cai, S. Osher, and Z. Shen. Linearized Bregman iterations for frame-based image deblurring. *SIAM J. Imaging Sci.*, 2(1):226–252, 2009.
- [32] X. Zhang, M. Burger, X. Bresson, and S. Osher. Bregmanized nonlocal regularization for deconvolution and sparse reconstruction. *SIAM Journal on Imaging Sciences*, 3:253–276, 2010.
- [33] S. Setzer. Split bregman algorithm, douglas-rachford splitting and frame shrinkage. *Scale space and variational methods in computer vision*, pages 464–476, 2009.
- [34] X. Zhang, M. Burger, and S. Osher. A unified primal-dual algorithm framework based on bregman iteration. *Journal of Scientific Computing*, 2010. .
- [35] M. R. Hestenes. Multiplier and gradient methods. *J. Optimization Theory Appl.*, 4:303–320, 1969. ISSN 0022-3239.
- [36] Dimitri P. Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Computer Science and Applied Mathematics. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York, 1982.
- [37] J. Eckstein and D.P. Bertsekas. On the douglas rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1):293–318, 1992. ISSN 0025-5610.
- [38] Nicholas J. Higham. *Functions of matrices*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2008. Theory and computation.
- [39] N.J. Higham. Newton’s method for the matrix square root. *Mathematics of Computation*, 46(174): 537–549, 1986. ISSN 0025-5718.
- [40] R.L. Burden and J.D. Faires. *Numerical Analysis*. 8th edition, Brooks/Cole, 2004.

- [41] J.F. Cai and S. Osher. Fast singular value thresholding without singular value decomposition. 2010.
- [42] N.J. Higham. Computing the polar decomposition- with applications. *SIAM J. Sci. Stat. Comput.*, 7 (4):1160–1174, 1986.
- [43] N.J. Higham and S.S. Robert. Fast polar decomposition of an arbitrary matrix. In *SIAM J. Sci. Stat. Comput.* Citeseer, 1990.
- [44] T.R. Hughes, M.J. Marton, A.R. Jones, C.J. Roberts, R. Stoughton, C.D. Armour, H.A. Bennett, E. Coffey, H. Dai, Y.D. He, et al. Functional discovery via a compendium of expression profiles. *Cell*, 102(1):109–126, 2000. ISSN 0092-8674.